

AD A 025000

Semiannual Technical Summary

Information Processing
Techniques Program

Volume I:
Packet Speech/Acoustic Convolvers


30 June 1975

Prepared for the Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-73-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Approved for public release; distribution unlimited.



The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Advanced Research Projects Agency of the Department of Defense under Air Force Contract F19628-73-C-0002 (ARPA Orders 2006 and 2929).

This report may be reproduced to satisfy needs of U.S. Government agencies.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency of the United States Government.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Eugene C. Raabe
Eugene C. Raabe, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

INFORMATION PROCESSING TECHNIQUES PROGRAM
VOLUME I: PACKET SPEECH/ACOUSTIC CONVOLVERS

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
ADVANCED RESEARCH PROJECTS AGENCY

1 JANUARY - 30 JUNE 1975

ISSUED 16 OCTOBER 1975

Approved for public release, distribution unlimited.

LEXINGTON

DDC
RECEIVED
JUN 2 1976
A
MASSACHUSETTS

ABSTRACT

The Information Processing Techniques Program sponsored by DARPA at Lincoln Laboratory consists of three efforts: Packet Speech, Acoustic Convolvers, and Airborne Command and Control. In this Semiannual Technical Summary, the first two areas are reported in Vol. I and the third in Vol. II.

ACCESSION for		
NTIS	White Section	<input checked="" type="checkbox"/>
DOC	Buff Section	<input type="checkbox"/>
UNANNOUNCED		<input type="checkbox"/>
JUSTIFICATION		
BY		
DISTRIBUTION/AVAILABILITY CODES		
Dist.	AVAIL. and/or SPECIAL	
A		

CONTENTS:

Abstract	iii
I. <u>PACKET SPEECH; <i>only</i></u>	1
A. Overall Aim of Packet Speech Program	1
B. Activities of the NSC Group	1
C. Summary of Lincoln's Activities in NSC	1
D. Network Measurements	2
E. Network Speech Experiments	2
F. Digital Voice Terminal (DVT)	3
G. Development of a Network-Speech Facility	4
H. Audio Pre-processing of Speech for NSC Use	5
I. CVSD Work	5
1. Conferencing Hardware	5
2. CVSDM-PLI Tests	7
J. Voice-Excited Linear Prediction Vocoder	7
1. Introduction	7
2. General System Description	9
3. Illustration of System Operation	10
4. Speech Quality Evaluation	10
K. Miscellaneous	13
II. <u>ACOUSTOELECTRIC CONVOLVERS.</u>	13
A. Introduction	13
B. Acoustoelectric Convolver	19
C. Convolver Test Set	21
D. Results	22
E. Conclusions	22
F. Coherent Integrators	22

INFORMATION PROCESSING TECHNIQUES PROGRAM

I. PACKET SPEECH

A. OVERALL AIM OF PACKET SPEECH PROGRAM

The long-range objective of the Packet Speech Program is to develop and demonstrate techniques for efficient digital speech communication on networks suitable for both voice and data, leading eventually to better interaction between people and computers in a multicomputer network environment. The ARPA Network has demonstrated the advantages of message-switched data networks. Presently, the ARPA-sponsored Network Speech Compression (NSC) group is studying the effects of the network on digital speech transmission. Ultimately, development of a successful voice-data network will require application of many of the results of the past half-century of speech research (drawing upon such topics as speech recognition, speech synthesis, speech bandwidth compression, and speaker authentication) for effective digital speech communication. Current efforts of the NSC group are strongly oriented toward adaptation of the ARPANET to fulfill voice needs and the problems of speech bandwidth compression.

B. ACTIVITIES OF THE NSC GROUP

A major joint activity of the NSC group has been the development of a packet speech capability. Of the eight NSC contractors, Lincoln, Information Sciences Institute (ISI), Stanford Research Institute (SRI), Speech Communication Research Laboratory (SCRL), and Bolt, Beranek and Newman (BBN) have procured, via ARPA, the SPS-41 computer for use as a real-time speech processor. Each of these contractors also has PDP-11/40 or PDP-11/45 facilities so that both speech algorithms and network protocol programs could be shared by all the above users. In May and June of 1975, several partially successful speech demonstrations were carried out between ISI and a facility at Lincoln composed of the fast digital processor (FDP) and the TX-2. To date, no reliable speech communication facility is available to NSC, pending resolution of difficulties encountered with processor hardware by ISI and SRI.

C. SUMMARY OF LINCOLN'S ACTIVITIES IN NSC

Lincoln has continued its work on ARPA network measurements. The results of these measurements should strongly influence the methods of efficiently processing speech packets to obtain good speech quality. Type 3 packets [raw packets sent without RFNMs (Request for Next Message)] suggested by Dr. R. E. Kahn have proved to yield appreciably higher bandwidth transmission. As an alternate to the SPS-41, Lincoln was asked by ARPA to construct two digital voice terminals (DVTs) for use as real-time speech processors. Four DVTs have already been constructed for other sponsors and two more DVTs will be constructed for ARPA during calendar 1975. Lincoln is also developing a network speech facility which can interface with either a DVT or an SPS-41. When this facility is operational, NSC members will be able to talk to each other using either DVTs, SPS-41's or the CHI (Culler-Harrison Incorporated) array processor.

In addition to the above, Lincoln has completed its work on VELP (voice-excited linear prediction), is working jointly with M.I.T. in developing a flexible audio front end, and has helped BBN check out their PLIs (private line interface) in a speech environment.

D. NETWORK MEASUREMENTS

During this reporting period, a new type of message service has become available in the ARPANET. This service makes use of Type 3 or "raw" packets and differs from regular (Type 0) message service in that the network assumes no end-to-end responsibility for successful delivery of messages. Packets are sent out from the source IMP (Interface Message Processor) without waiting for any indication from the destination IMP to the effect that buffer space will be available for the packets when they arrive. If such space is not available, Type 3 packets are discarded on arrival with no indication of such action being given to either the sending or receiving host. Further, Type 3 packets are delivered to the receiving host in the order of their arrival at the destination IMP, which is not necessarily the order in which they were given to the source IMP by the sender.

The advantages of Type 3 messages over Type 0 were expected to be greater data-rate capability and reduced transmission time through the network. The penalties to the user were expected to be some possibility of lost messages and increased overhead concerned with reordering the order of received messages.

The bulk of our measurements work this period has been concerned with Type 3 messages and the comparison of their performance with respect to Type 0 messages. Because there is no RFNM returned to the sending host for Type 3 messages, the "fake-host 3" measurement technique described in the previous SATS could not be used; instead, a new measurement technique was developed. This new technique requires the cooperation of a remote host in running a program called ECHO which immediately retransmits received messages adding a time stamp and a sequence number indicating the time and order in which the message was received.

Using an ECHO program running in a PDP-11 at ISI, we have made a series of measurements of network delays for a variety of data rates. The results are discussed in detail in NSC Note 70, "ARPANET Delay Measurements," which has been distributed to the NSC community. The data from the measurements confirm the expected superiority of Type 3 messages for speech communication purposes. Delays observed with Type 3 messages were less than those observed with Type 0 messages throughout the range of data rates we explored (3.5 to 18.5 kbps). No Type 3 messages were lost, and out-of-order arrival was appreciable only at the higher data rates (6.8 percent at 18.5 kbps to 151). During the period when the measurements were taken, we were essentially the only active users of the Type 3 message capability; network performance might be different as regards number of messages lost if Type 3 messages were being used by many hosts on the net at the same time.

E. NETWORK SPEECH EXPERIMENTS

The TX-2/FDP program for LPC (linear predictive coding) vocoded speech communication on the ARPANET was modified to handle the Type 3 messages discussed above in connection with network measurements. The modifications provided for the proper ordering of speech packets which might arrive out of order. New statistics-gathering capabilities were added to allow the measurement of the effectiveness of the playout algorithms in smoothing the jitter in message delays through the net. Further modifications allowed the use of the ECHO program at a remote host to send vocoded speech back to TX-2 so that we could hear our own transmitted speech and thereby simplify the debugging of code as well as explore the effectiveness of playout algorithms.

During the period, a number of successful network conversations were held with both ISI and CHI. Using the ECHO program at ISI, we produced an analog tape of a conversation with about 1.4-sec round-trip delay including smoothing delays in the playout algorithm. With the delay settings used for this tape, about 1.5 percent of the speech packets arrived too late to be used by the playout algorithms and consequently resulted in "glitches" in the output speech. In our judgment this glitch rate, while detectable, is not too disturbing and represents a reasonable target when seeking a compromise between round-trip delay and glitch rate.

Network speech experiments, which had formerly used the TX-2 facility, aimed at improving the playout algorithm will be resumed following implementation of network speech programs on our PDP-11/SPS-41 system. The important decision which that algorithm must handle is when to begin playing out speech after a silent interval. If playout is started too soon, a glitch will result if the playout catches up with the flow of data from the net. If playout is delayed too much, unnecessary delay will be introduced in the conversation. The algorithm used in the early versions of our program (also in those at ISI and CHI) simply waited a fixed time to start playing out after the arrival of the first speech packet following a silent interval. In effect, that first packet was used as an estimate of the mean network delay and the fixed time delay was added to allow for the variance in network delay. Unfortunately, that first packet is not a good estimate of the mean, and a relatively large fixed delay must be used to avoid excessive glitching. Our more recent programs have used a running average of differences between arrival times and a local clock to adjust the clock to a better approximation to the mean network delay characteristic. Preliminary results show that the overall delay can be reduced with such an algorithm without increasing the glitch rate.

F. DIGITAL VOICE TERMINAL (DVT)

Under Air Force sponsorship, Lincoln has designed and constructed a digital processor capable of performing a variety of speech bandwidth compression algorithms. The DVT is a programmable processor capable of executing 18 million instructions/sec. It consists of 4 wire-wrapped boards containing a total of 470 ECL (emitter-coupled logic) packages and 125 TTL (transistor-transistor logic) packages. The DVT fits into a 19- x 5- x 25-inch frame. This structure contains its own power supply and is thus self contained and can be connected directly to a modem (or a PDP-11) for digital speech transmission.

At the present writing, the following speech algorithms* have successfully run on the DVT: the LPC algorithm at 2400, 3600, and 4800 bits per second (bps) rates; the APC (adaptive predictive coding) algorithm at 8000 bps; two different TRIVOC (triple-function vocoder) algorithms, both at 2400 and 3600 bps; and the ARC (adaptive residual coding) algorithm at 9600 and 16,000 bps. These algorithms have been developed over a period of several months. By contrast, the single LPC algorithm developed by NSC for the SPS-41 has taken more than one year. The primary reasons for this are (1) hardware problems in the SPS-41 and (2) much greater ease of programming the DVT. In this latter connection, it is important to note that the DVT is a highly sequential machine whereas the SPS-41 is a parallel processor. By taking advantage of increased circuit speeds, a 55-nsec instruction cycle time became feasible and it was determined that parallel processing techniques were no longer needed for real-time speech processing.

*E. M. Hofstetter, P. E. Blankenship, M. L. Malpass, and S. Seneff, "Vocoder Implementations on the Digital Voice Terminal," 1975 EASCON Proceedings.

A future possible advantage of employing DVTs for packet speech work on the ARPANET is the ability of the DVT to rapidly switch from one program (say LPC) at a given data rate (say 3600 bps) to another program such as ARC at 9.6 or 16 kbps. This flexibility of bit rate with resultant variation in speech quality seems eminently suitable for a packet speech network where speech bandwidth transmission capabilities might fluctuate with the time-varying traffic.

Under ARPA sponsorship, two DVTs are presently being constructed and are scheduled to be completed in mid-October 1975 (DVT 5) and early January 1976 (DVT 6). During this construction phase, packet speech experiments can be conducted using some subset of the four presently existing DVTs, given the permission of the sponsors of these devices.

G. DEVELOPMENT OF A NETWORK-SPEECH FACILITY

Lincoln participated in the initial network-speech experiments with the TX-2 and FDP providing the computational resources. One aim of NSC is the establishment of nearly identical computer facilities for network-speech communication at each site. This common facility consists of an SPS-41 to process the speech signals and a PDP-11 interfaced to the ARPANET to handle the flow of messages through the network. Because of this standardization of facilities and the phase-out of the TX-2, Lincoln is installing such a facility.

The installation of the hardware is nearly complete. A PDP-11/45 with a 64K memory was installed last fall and its interface to the Lincoln IMP was completed last March. An SPS-41 was installed last December and the engineering changes required for the network-speech programs were made by SPS last May. Once A/D and D/A converters are installed on the SPS-41, the hardware installation of Lincoln's network-speech facility will be completed.

While our computational facility is nearly identical to those of ISI, SRL, SCRL, and BBN, the installation of the software necessary for network-speech communication proved to be time-consuming. The ELF operating system provided by SCRL was chosen as the medium through which network-speech communications would be made. A sophisticated version of ELF (VM or virtual-memory ELF) is required for the network-speech application. Toward installing this version of ELF, a more primitive, non-VM, version has been successfully generated and installed on the PDP-11.

As ELF lacks components necessary for normal operation by everyday users, an alternate operating system for the PDP-11 is required for daily use. Lincoln chose Digital Equipment Corporation's RSX-11M as its general-purpose operating system. System programs for assembling and linking the network software had to be modified for operation under RSX-11M. SPS-41 system programs - an assembler and a debugging program SPUD - were originally written for compatibility with the DOS operating system. A program (ARF) necessary for the creation of the overlay streams used by the LPC programs was written for operation on a PDP-10. These programs have all been successfully modified to run under RSX-11M.

Once these modifications were made, the SPS LPC programs written by ISI and SRL were installed. ISI's modifications to SPUD, which load these LPC programs, also have been installed. Consequently, once A/D and D/A converters are installed on the SPS-41 and a system for processing network messages is obtained, Lincoln will be ready to return to active participation in network-speech experiments, using either the DVT or the SPS-41 as the speech processor.

H. AUDIO PRE-PROCESSING OF SPEECH FOR NSC USE

At an ARPA NSC meeting held in Santa Barbara in December 1974, Lincoln agreed to work on the standardization problem for speech input to a digital processor. This pre-processing includes the choice of microphone, pre-emphasis filter, gain control, pre-sampling filter, de-emphasis filter and post-sampling filter.

Since an important ARPA long-range aim is the development of remote voice-data interactive facilities, it seems appropriate to use speech inputs which are close approximations to commercial equipment. To this end, Lincoln proposes that present speech work be done with standard telephone handsets with the carbon button microphone replaced by dynamic cartridges. Lincoln has supplied other NSC members with this equipment. In addition, a model audio front end has been designed at Lincoln which includes a wide variety of options for pluggable modules for the various devices cited in the above paragraph; this subject is covered in some detail in NSC Note 57.

It has been the experience of both NSC and other speech bandwidth compression workers that a wide variety of sampling rates are desired. Each time the sampling rate changes, the pre-sampling and post-sampling filter cutoff frequencies also should be changed. One way of accomplishing this is by plugging in a different analog filter each time. However, this means that any demand for a new sampling rate requires the construction of a new analog filter. As a step toward solving this difficulty, M.I.T. and Lincoln are collaborating on the construction of a programmable digital pre-sampling filter. The design of this filter is being carried out by M.I.T. and the construction will be carried out by Lincoln.

I. CVSD WORK

During this report period, the CVSDM (continuously variable slope delta modulation) waveform encoding equipment has been used for two ARPA applications. One application uses six CVSDM encoder-decoders driving a single PDP-11/45 interface structure for switchboard and conferencing experiments on speech data streams from six users. A second application has provided CVSDM encoder data streams for use with PLI bit stream devices.

1. Conferencing Hardware

Upon delivery of the CVSDM units (the end of calendar 1974), six were dispatched to the NSC group at ISI. During the spring of this year, D. Cohen's group at ISI developed a combiner device designated the PB11 which accepts data in from and out to several CVSDM devices, combines them into a single parallel handshake for the PDP-11/45, and provides each device with a common synchronized clock. Using a PB11, it will be possible to combine six CVSDMs (or more) into a single PDP-11/45 as a switchboard and use this common node as the basis for multi-user voice access to the net in both a single-user extension mode and a many-user conference mode. The plan is to reproduce a PB11 box at Lincoln so that the Lincoln PDP-11/45 can communicate with several CVSDM devices connected to it and also communicate with ISI, through the network, in several conferencing modes at CVSDM rates.

2. CVSDM-PLI Tests

In mid-March 1975, two CVSDM units were lent to BBN for use as speech input devices for PLI development and testing. The PLI, insofar as packet speech is concerned, looks like

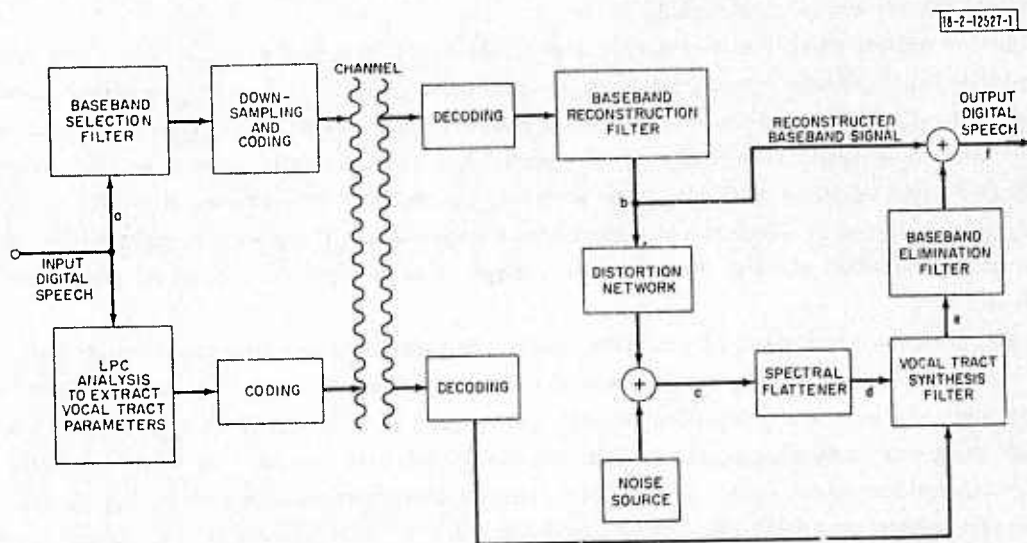


Fig.I-1. Block diagram of voice-excited linear prediction (VELP) vocoder.

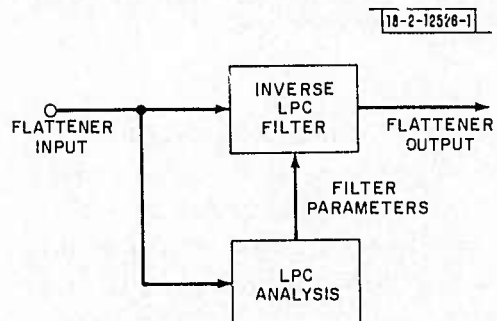


Fig.I-2. Spectral flattening technique in VELP.

a minihost capable of accepting and delivering serial speech data while providing properly formatted packets to a network IMP. Working with BBN personnel, Lincoln staff provided guidance for proper PLI speech software development including buffer delay for net smoothing and silence detection strategy.

Using this software, PLI-CVSDM experiments were performed by BBN and voice transmitted through various network connections. A recording of speech input and output from experiments on from one to ten link paths was made using various CVSDM rates. As expected, the longer hops could not sustain the 16-kbps rate using Type 1 message format, although the shorter hops sound fine (i.e., small delay and no discontinuities). This recording is on file with Dr. R. E. Kahn.

J. VOICE-EXCITED LINEAR PREDICTION VOCODER

1. Introduction

The previous SATS included a preliminary description of a voice-excited linear prediction (VELP) vocoder with adaptive spectral flattening. Further work on this vocoding scheme has resulted in an 8-kbps VELP system with output speech quality roughly comparable to that of an 8-kbps APC system.* In addition, VELP quality seems less degraded than APC, or than pitch-excited LPC, when the input speech is of telephone quality. A condensed description of the current VELP system follows; a more detailed description has been prepared for publication.†

2. General System Description

An overall block diagram of the VELP system is shown in Fig. 1-1. As in a voice-excited channel vocoder, the excitation information is transmitted in an unprocessed (except for sampling and quantization) sub-band of the original speech. The baseband limits must be chosen so that at a minimum, either the voice fundamental, or else the second and third harmonics of the speech, are transmitted. With this requirement satisfied, the nonlinear distortion network is able to generate a waveform containing all the original harmonics of voiced input speech. The spectral shape superposed on this regenerated harmonic structure is quite arbitrary and must be flattened to produce a suitable excitation function. In the channel vocoder case, this was accomplished by a spectral flattener consisting of a bank of bandpass filters, each followed by a limiter. Spectral flattening in VELP is achieved by carrying out an LPC analysis of the distorted baseband signal, and passing this signal through an inverse filter based on the derived predictor coefficients, as shown in Fig. 1-2.

The description of excitation function generation has so far been phrased in terms of voiced sounds, where the problem is most difficult. To account for unvoiced sounds, a small amount of white noise is added to the distorted baseband. This noise will be essentially unnoticed during voiced sounds, but will produce the desired hiss excitation for unvoiced sounds. In order that the excitation energy be correct for both voiced and unvoiced sounds, a gain parameter is included in the flattening filter so that its output energy is equal (on a frame-by-frame basis) to the energy in the residual error signal of the original speech.

The excitation signal is passed through an LPC synthesis filter whose parameters have been determined at the transmitter by LPC analysis of the original speech. The final step is

*E. M. Hofstetter, P. E. Blankenship, M. L. Malpass, and S. Seneff, "Vocoder Implementations on the Digital Voice Terminal," 1975 EASCON Proceedings.

†C. J. Weinstein, "A Linear Prediction Vocoder with Voice Excitation," 1975 EASCON Proceedings.




Fig.I-3. Spectrograms of original speech (top) and VLP output (bottom). The utterance is "We've changed the measures," spoken by a male talker.

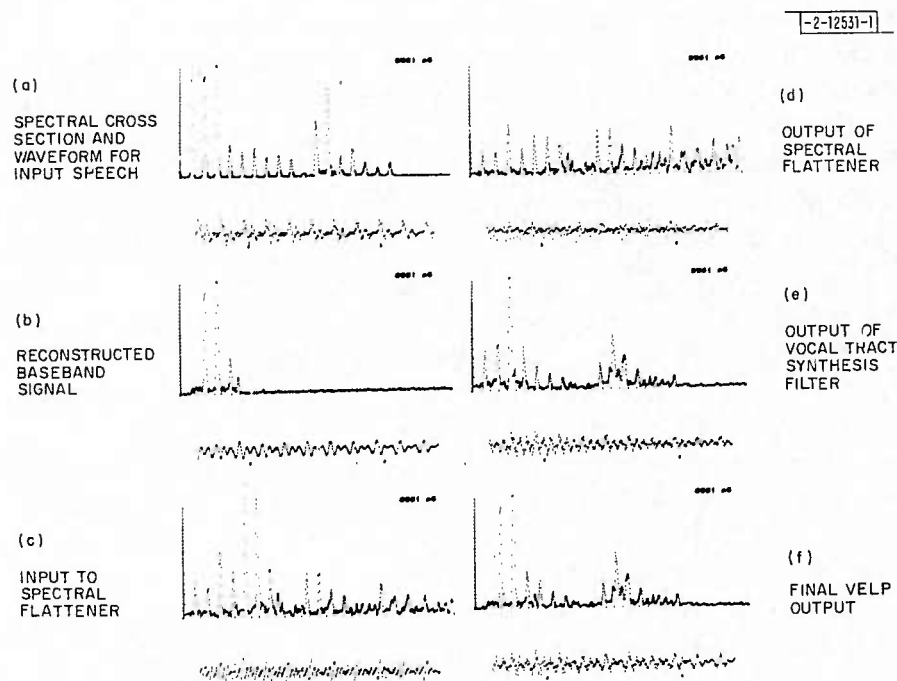


Fig.I-4. Spectral cross sections and associated waveforms at various points in VLP system. The spectral cross-section plots represent spectral amplitude vs frequency, with the frequency axis covering 0 to 3572 Hz.

to filter out the baseband part of this synthesized speech and then add back in the essentially clear band of original speech which was transmitted over the channel. This direct inclusion of the transmitted baseband signal in the output speech is an important positive feature of voice-excited vocoders in general, and VELP in particular.

3. Illustration of System Operation

We will now illustrate VELP operation by means of spectrograms, spectral cross sections, and waveform segments at various points in the system. The system used in obtaining this data has an overall bit rate of 8 kbps, LPC of 8th order both for vocal tract analysis and for flattening, and a distortion network consisting of an asymmetrical full-wave rectifier followed by a differencer. Complex sampling is used on a baseband extending approximately from 150 to 750 Hz, and coding of the baseband is 5-bit log PCM. The actual rate for the baseband is 5950 bps, representing 595 complex samples/sec. With the spectral data, the precise overall bit rate is thus 8072 bps.

Wideband spectrograms of the input speech and VELP output for the utterance, "We've changed the measures," spoken by a male talker, are shown in Fig. I-3. It is seen that the general spectral shape is fairly well maintained. Probably the most visible effects of VELP are the irregularity of the pitch striations and the generally noisy appearance of the output speech spectrum. These effects are perceived as a slight hoarseness in the output speech. The explanation seems to be that the process of regenerating and flattening the speech harmonics leads to an excitation spectrum which is not as clean as would be desired, in that the spectral peaks representing the harmonics are broadened, and the ratio of the level of the harmonics to the spectral hash between harmonics is lowered. This effect, and other aspects of system operation are more clearly seen in the spectral cross section pictures of Figs. I-4(a-f), which are taken during the vowel in the word "changed," as indicated by the markers under the spectrograms.

In each of the Figs. I-4(a-f), a spectral cross section is shown above the waveform segment analyzed. The pictures represent measurements at points labeled a through f in Fig. I-1. Each spectrum analysis was carried out at the same location in time with respect to the beginning of a digitally stored version of the input utterance. The sampling interval was 140 μ sec, and spectral amplitude values (linear scale) are displayed from 0 Hz to the half sampling frequency, 3572 Hz. Each spectrum was calculated by a DFT (discrete Fourier transform) of 512 samples of hanning-windowed waveform. The spectral samples produced by the DFT are therefore spaced by $7144/512 = 14$ Hz. Every fourth point on the spectrum displays represents an actual DFT magnitude output. The remaining points were filled in by linear interpolation. The spectral analysis parameters were selected as described in order to provide a revealing display of the harmonic structures of the various spectra.

Figure I-4(a) represents the digital input speech, produced by low-pass filtering, pre-emphasis, and sampling of the analog speech. The harmonic structure as well as the spectral envelope shape (characteristic of a nasalized vowel) are apparent. Figure I-4(b) represents the baseband signal as reconstructed at the receiver. The 2nd, 3rd, and 4th harmonics have been retained with fairly good fidelity but the combined effects of aliasing due to filter side lobes, PCM coding noise, and filter roundoff errors have introduced some small, undesired signal components both inside and outside the baseband. The input to the spectral flattener is shown in Fig. I-4(c). Many of the harmonics have been regenerated, but the amplitude envelope of the

harmonics is far from flat. It is apparent also that the harmonic structure is not nearly as clean as in the original speech. Many of the low-level harmonics are broadened and smeared, some harmonics appear to be missing, and a visible hash level is superposed on the entire spectrum. The output of the spectral flattener is shown in Fig. I-4(d). The larger variations in spectral envelope have been removed, although sizable local variation among adjacent harmonics remains. Such local variations tend to cause important problems only if they occur in the vicinity of a sharp vocal tract resonance. In this example, and in many other cases which have been observed, the worst combination of lack of excitation flatness plus strong vocal tract resonance occurs in the baseband itself. Thus, if we compare the synthesis filter output spectrum in Fig. I-4(e) to the original spectrum of Fig. I-4(a), we notice large spectral distortion in the baseband. For example, the 3rd harmonic is more than twice as large as the 2nd in Fig. I-4(e). Fortunately, this effect, which is quite important perceptually, is corrected by the re-introduction of the baseband signal into the output speech. In the final output shown in Fig. I-4(f), we see that the amplitude relationship of the 2nd and 3rd harmonics is the same as in the original speech. Of course, the fundamental has been removed by the baseband filter and is missing in the output speech, since the baseband elimination filter employed in this system was high-pass at 750 Hz.

4. Speech Quality Evaluation

The following comments on VELP are based on informal listening tests and reflect the opinions of several experienced listeners. Detailed subjective testing of VELP has not been carried out, but fairly extensive listening was possible, since the system runs in real time on the FDP at Lincoln.

The 8-kbps VELP system was compared against an 8-kbps APC system and a 3.5-kbps pitch-excited LPC system, each implemented on the FDP. For high-quality speech inputs, VELP and APC produced output speech of roughly similar quality. However, a small residual hoarseness in VELP caused most listeners to express a preference for APC. The LPC system possessed the usual lack of naturalness associated with a pitch-excited system. The same three systems were compared on a recorded sample of speech which had been transmitted through a local dial-up connection in the public telephone system. Here the results were somewhat different, as pitch glitches became somewhat noticeable in APC and extremely bothersome in LPC. The VELP quality was least degraded by the telephone speech, supporting the assertion that voice-excitation is relatively robust in this type of environment.

K. MISCELLANEOUS

Our work on variable-rate coding algorithms for LPC has presently come to a halt. In the previous SATS, we showed results of our equal-area coding scheme. We made an attempt to compare the efficiency of this scheme with that of the more conventional log-area coding scheme used by NSC. Our criterion of comparison was the mean-squared difference between the uncoded and coded reflection coefficients for the two schemes. No definitive results could be obtained; on further reflection we concluded that only a comparison based on subjective speech tests would be significant. What remains to be done for the speech compression community at large is a comprehensive test program aimed at determining the best of a variety of proposed coding schemes.

In our judgment, LPC speech quality varies quite noticeably among speakers. In particular, female speakers appear to cause greater difficulty than male speakers. In a preliminary attempt to understand this problem, we have modified an existing real-time LPC program so that a given speaker of category male, female, or child can be perceptually transformed into a different category. Based on informal listener responses, the results have been quite successful. To perform this transformation we adjust two parameters independently: the pitch and the spectrum. The pitch adjustment is easily carried out by multiplying the measured pitch by a scale factor while the spectral adjustment involves an apparent time speedup (such as would be attained by changing the playback speed of a tape recorder) but maintaining the overall speech rate output the same as the input rate. If a simple nondistortionless method of implementing this transformation could be found, it could result in a significant improvement in LPC quality.

II. ACOUSTOELECTRIC CONVOLVERS

A. INTRODUCTION

The Lincoln Laboratory effort on acoustoelectric convolvers for packet radios is motivated by a desire to increase the instantaneous bandwidth and signal-processing gain of continuously variable, pseudorandom-code modems. Collins Radio is currently building for ARPA several packet radio repeaters. Lincoln Laboratory is to supply the first three acoustoelectric convolver subsystems for the repeaters.

The convolver subsystem is being developed in two steps. An initial convolver prototype was developed during this report period which meets the requirements for a single-bit convolver, and two convolvers and a test circuit were built and evaluated at Lincoln Laboratory. The finished units are currently being transferred to Collins Radio for their evaluation. We are converging on a final specification for a convolver subsystem for the packet radio repeaters, and the current subsystem specifications are contained in this report. Also, the acoustoelectric convolver technology is being transferred as rapidly as possible to industrial sources.

We are continuing the study of rapidly acquiring a time-slotted signal, with a continuously changing code structure. Our current plans are to develop a circuit which will permit the rapid acquisition of an incoming differential-phase-shift-keyed (DPSK) signal.

We are exploring the feasibility of a coherent integrator, a new device concept which might provide signal processing gains as large as 10^5 . The coherent integrator is an outgrowth of our convolver technology. It consists of a block of lithium niobate with a strip of silicon which is spaced a fraction of a micrometer away from the lithium niobate surface. The silicon interface is covered with a two-dimensional array of Schottky diodes with polysilicon overlays. Some initial results are included in this report.

B. ACOUSTOELECTRIC CONVOLVER

Two views of the finished convolver package are shown in Fig. II-1. The basic assembly procedures are similar to those described in the previous SATS.* The silicon strip, which is 30 mils wide and $1\frac{1}{2}$ inches long, has a backing of metal, and 10-mil gold wires are bonded to the middle and ends of this strip. The silicon rests on, and the wires pass through a block of silicone potting gel (RTV-602). Just before assembly is attempted, the lithium niobate and silicon interfaces are coated with liquid collodion. The liquid dries rapidly, and any dust on the surfaces is captured in the congealed film. The film is pulled off the surfaces just prior to assembly and the freshly cleaned surfaces are brought into contact. The brass mounting structures are bolted together, which compresses the RTV sufficiently to apply 0.1-atm pressure to the silicon relative to the lithium niobate spacer-posts.

The input signals are at 300 MHz, and when two 10- μ sec-long signals convolve beneath the silicon, then the output signal is at 600 MHz. Since the dielectric constant of LiNbO_3 is approximately 50, the 1.4-inch silicon acts like a half-wavelength-long transmission line. The metal ground planes, which are visible on the surface of the lithium niobate [Fig. II-1(a)] are designed to provide a characteristic impedance of 30 ohms to this Si- LiNbO_3 transmission line. A considerable amount of distortion would be obtained, due to the long-line effects in the output

*Information Processing Techniques Program Semiannual Technical Summary, Lincoln Laboratory, M.I.T., Vol. I (31 December 1974).




Fig. II-1(a). Open convolver structure showing lithium niobate and silicon strip. The lithium niobate has deposited on its surface metallic transducers and ground planes. The silicon strip is on a pillow of RTV gel. The silicon is pressed against the stand-off spacers located in the gap between ground planes.




Fig. II-1(b). Convolver with cover removed. The inner brass block contains the output transformer in a central cavity and two terminations for the silicon transmission line are in two cavities at the edge. An airtight seal is formed between the outer block and the cover, and the device is coated with a sealant. The package housing is evacuated and back-filled with dry nitrogen through the entry hole located adjacent to the output terminal. Quarter-wave impedance inverter networks are attached to the two input terminals.

circuitry, unless care is taken to terminate the ends of the silicon with the characteristic impedance of the 30 ohms. That is why a wire is bonded to each end of the silicon. The bonding wire has a 10-ohm inductive reactance which is compensated with a 26-pF capacitor. The output terminal at the center of the silicon has a characteristic impedance of 15 ohms, which is connected to a 4:1 impedance transformer. In Fig. II-1(b) the output transformer is visible in the center of the inner brass block, as well as the terminations connected to each end of the silicon strip. The input transducers have a wide fractional bandwidth of 30 percent, which results in an input impedance of about 300 ohms of reactance in series with 50 ohms of radiation resistance. A quarter-wave impedance inverter network is attached to the input terminal. It consists of a quarter-wavelength-long cable with a characteristic impedance of 50 ohms, a two-to-one impedance transformer, and tunable inductor [see Fig. II-1(b)]. The input impedance over the frequency band of interest is shown in Fig. II-2(a). Notice that the input impedance follows a constant K circle on the Smith chart, which is equivalent to a VSWR of 3.0, and a reflection loss of 1 dB. The output impedance plot is shown in Fig. II-2(b). The residual reactance at the output is in part due to stray reactance associated with the output transformer. Note that the VSWR is 3, as in the input circuit.

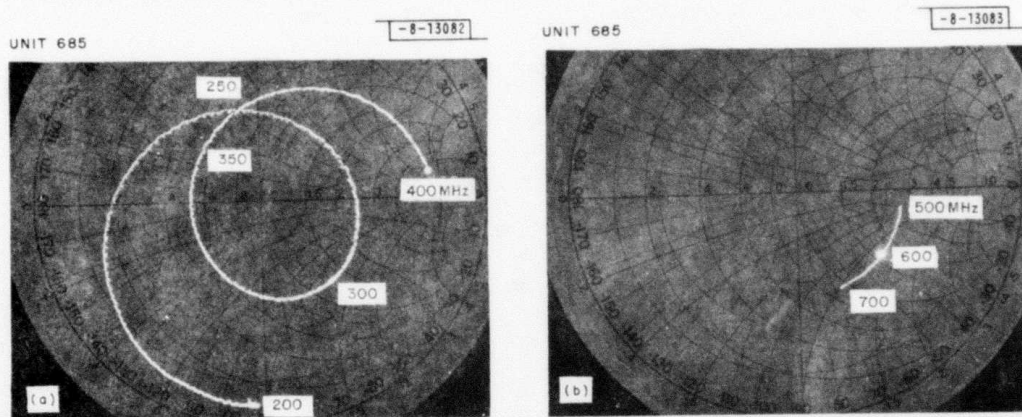


Fig. II-2. (a) Smith chart plot of input impedance. Frequency of interest is from 250 to 350 MHz. Effective VSWR is 3:1. (b) Smith chart plot of the output terminal impedance from 500 to 700 MHz. Effective VSWR is approximately 3:1.

Four convolvers were built on substrates 675, 680, 682, and 685. Three of the convolvers have essentially identical characteristics. The fourth convolver, unit 682, has an excess insertion loss of 10 dB, which has been traced to a residue of collodion particles on the crystal face. A critical test of the uniformity of convolvers is to convolve a short CW impulse with a long pulse. If the convolver interaction is uniform, then the expected output signal should remain constant as the short impulse traverses the convolver interaction region. Figure II-3 shows this output for unit 685. A variation of ± 1 dB is obtained in this device. This variation was traceable to a residual impedance mismatched at the ends of the silicon. Unit 680, with improved terminations, does not exhibit such fluctuations.

The convolution of two CW pulses 10 μ sec long creates a characteristic triangular waveform. The peak amplitude of this convolution signal as a function of frequency is plotted in Fig. II-4 as




Fig. II-3. Convolution of a 1.3- μ sec impulse with a 24- μ sec pulse. Notice the variation of convolution signal as a function of position. This is due in part to imperfectly terminated silicon strips.




Fig. II-4. Output signal levels as a function of input frequency for 10- μ sec CW pulses at 10 dBm. Notice that the signal is uniform to ± 0.5 dB over the 100-MHz bandwidth.





Fig. II-5. Output signal level at 600 MHz in dBm vs the product of the 300-MHz input signals in dBm. The two input signals are kept equal and their logarithms are summed in this figure. The 1-dB compression point occurs when both input signals are at 14 dBm.

a function of input frequency. Notice that the response of the convolver is flat within $\pm 1\frac{1}{2}$ dB over the 100-MHz band of interest. The convolver is a nonlinear element in the sense that the output signal is proportional to the product of the input signals. For that reason, the output amplitude is plotted as a function of the sum of the log (in dBm) of the two input signals in Fig. II-5. In this instance, the two inputs are kept equal and are varied in step. Notice that the output is linear with respect to the inputs up to input levels of +14 dBm, where a deviation from linearity of 1 dB is obtained. The output bandwidth is 200 MHz, and a reasonable noise floor created by an amplifier with a noise figure of 3 dB is -88 dBm. Consequently, the available dynamic range between this noise floor and the 1-dB compression point is 50 dB.

The principal spurious signal in the convolver is due to the self-convolution of an input signal. If a CW signal is entered at a terminal, then a reflection which occurs at the second terminal subsequently convolves with the input. The peak amplitude of this self-convolution signal is plotted in Fig. II-6 as a function of input power level. Notice that this self-convolution signal decays approximately as the square of the input power. The self-convolution of a coded signal is substantially less than the CW signal as a result of the lack of correlation between the beginning and the end of a code sequence. We have found with pseudorandom codes that the self-convolution signal is about 20 dB weaker than the self-convolution of a CW signal. Only two data points could be observed because of the excessive noise floor in our measurement equipment. Another spurious signal, the direct feed-through of the input signal to the output terminal, is 70 dB below the input. However, this is suppressed readily with a high-pass filter. No other spurious signals were observed with the convolver.

Fig. II-6. Spurious output signal level due to double-transit reflection vs input power level. Spurious signal decreases as the square of the input. Coded signals do not correlate, consequently they produce spurious signals which are 20 dB weaker than CW signals.



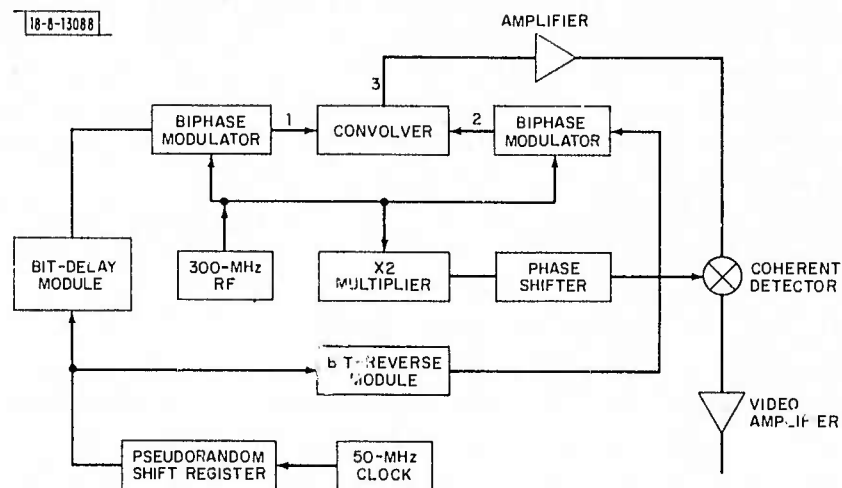


Fig. II-7. Acoustoelectric convolver test set schematic.

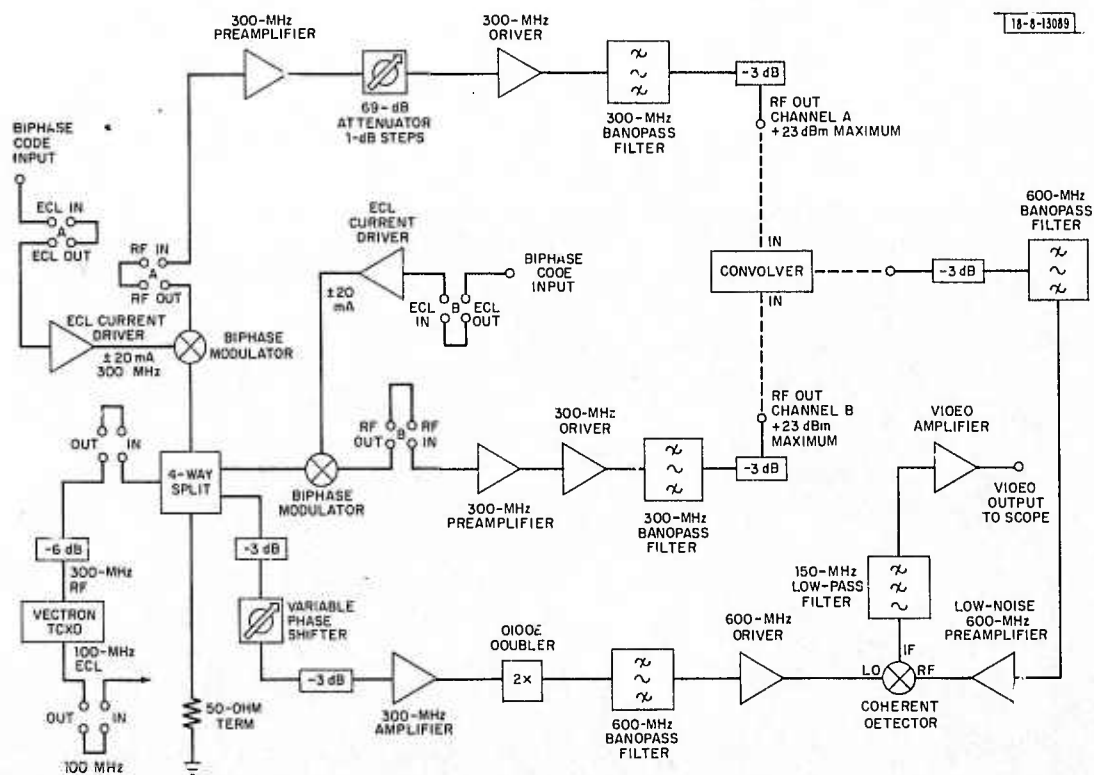


Fig. II-8. Block diagram of acoustoelectric convolver test set, showing electronic functions.

The performance characteristics of the acoustoelectric convolvers are summarized below. They conform to the requirements for the packet radio repeater, with the exception of VSWR, where the requirements are 2:1. A modest effort should correct this difficulty.

Center frequency	300 MHz
Bandwidth	100 MHz
Convolution interval	10 μ sec
Input VSWR	3:1
Output VSWR	3:1
Maximum reference power	+10 dBm
Maximum signal power	+10 dBm
Maximum output signal	-45 dBm
Convolution loss*	-66 dBm
Maximum double-transit signal	-70 dBm (CW); -90 dBm (coded)

C. CONVOLVER TEST SET

An acoustoelectric convolver test set was designed and built by Lincoln Laboratory, for use by the Collins Radio Company.

Figure 11-7 is a simplified system block diagram. A feedback shift register generates a pseudorandom code. This code is reversed through the bit reverser module. To correct for the time delay, the non-reversed code is sent through a similar bit-delay module. These pseudorandom codes modulate the 300-MHz RF in biphase modulators. The biphase modulated (0 to 180°) signals are applied to ports 1 and 2 of the convolver.

The two biphase coded RF signals propagate past each other in the convolver. Since one signal is a time-reversed replica of the other, an autocorrelation impulse is obtained in the convolver. Also, since both signals are propagating in opposite directions, the cluster frequency of the output is twice the input frequency, or 600 MHz.

The 600-MHz output is amplified and applied to a doubly balanced mixer used as a coherent detector. The local oscillator (LO) comes from the 300-MHz source which is doubled to 600 MHz and fed through a phase shifter to peak the detected response of the coherent detector. The resulting video signal is amplified and displayed on a wideband oscilloscope.

The acoustoelectric convolver test set is shown in more detail in the block diagram in Fig. 11-8. All frequencies are generated by a temperature-compensated crystal oscillator (TCXO). A 100-MHz, emitter-coupled logic (ECL) compatible signal is sent to the logic circuits where it is used as the digital system clock. A second TCXO output at 300 MHz is fed to a 4-port splitter. Two of the outputs are modulated by current drivers. Channel A is preamplified and sent through a step attenuator which is used to adjust the level of that channel. The attenuator output is amplified by a driver amplifier, filtered by a 300-MHz bandpass filter, and applied to output port A.

Channel B is fed straight through the preamplifier, driver, and filter to output port B. Maximum power of +23 dBm is available at output port B; and since it is run at a fixed level, it can be considered the reference channel.

A third output from the splitter is phase shifted, amplified, and doubled. The resulting 600-MHz signal is filtered and amplified. This signal is used as the LO of the coherent detector. The phase shifter is continuously variable to allow peaking of the coherent detector.

*Convolution loss = $10 (\log P_{\text{ref}} + \log P_{\text{sig}} - \log P_{\text{out}})$.

The fourth output of the power splitter is used for diagnostics and troubleshooting and is normally terminated in 50 ohms.

The return from the convolver is filtered by a 600-MHz bandpass filter and amplified by a low-noise preamplifier. The output of the preamplifier is applied to the coherent detector where it is mixed with a 600-MHz LO. The resulting difference frequency (a video pulse) is filtered by a low-pass filter and amplified by the video amplifier. The output is displayed on a wideband oscilloscope.

The front panel is illustrated in Fig. II-9. Channel A RF Out and its step attenuator are to the right, and Channel B RF Out is to the left. These two ports are hooked to the convolver under test. The 600-MHz return from the convolver is fed into the RF In port. The LO phase is adjusted with the ϕ adjust knob for maximum output on the display once the system is functioning.




Fig. II-9. Front panel of acoustoelectric convolver test set.

Several toggle switches control the digital logic board. The start/stop switch allows either a continuously changing pseudorandom code (START) or a fixed code (STOP) to be selected. The bit code and the corresponding switches allow the user to select the code length according to the following chart.

Switch 1 2 4				
0	0	0	8-bit code (2^8-1)	0 = down
1	0	0	11-bit code ($2^{11}-1$)	1 = up
0	1	0	33-bit code ($2^{33}-1$)	
1	1	0	36-bit code ($2^{36}-1$)	
0	0	1	N/A	
1	0	1	N/A	
0	1	1	N/A	
1	1	1	20-nsec square wave — for diagnostics	

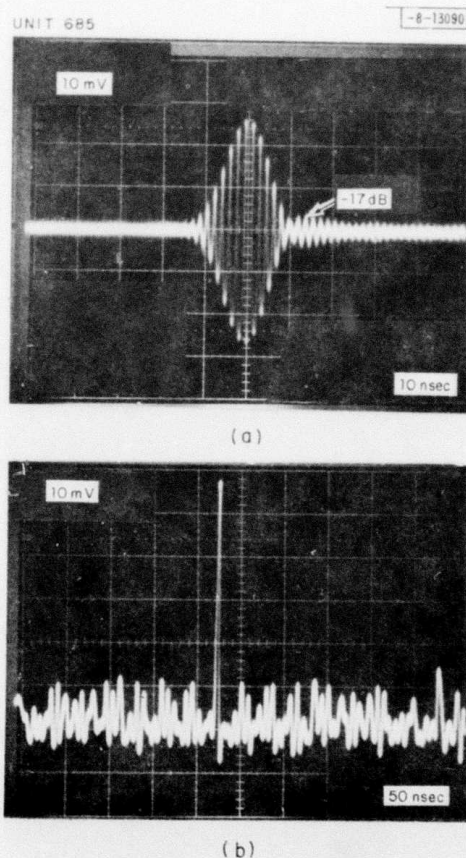
The ϕ reverse switch allows the biphasic code to be normal or inverted, thus changing the polarity of the autocorrelation pulse. The code generator reset allows the digital circuits to be reset in the event that all "0's" are entered in the code generator during turn-on or setup.

The acoustoelectric convolver test set has been designed to provide considerable versatility and flexibility. It should aid in the testing and analysis of the acoustoelectric convolvers currently under development by Lincoln for packet radios.

D. RESULTS

At present, a considerable amount of distortion is observed in the phase-modulated signal from the test set, a typical code sequence contains chips which are not the requisite 20 nsec long, and a considerable amount of amplitude modulation is present as well. These distortions tend to produce energy in a sidelobe following the main correlation impulse, which is on the average -17 dB below the main impulse. This can readily be seen in Fig. II-10(a) which shows the correlation impulse for a running code, the spurious sidelobe, while other sidelobes are blurred by the code. A typical sidelobe structure for a repeated code is shown in Fig. II-10(b) with time sidelobes as large as -14 dB below the main impulse.

Fig. II-10. (a) Correlation impulse obtained with a continuously changing RF code. Notice the correlation impulse has the expected 29-nsec duration time, and the time sidelobes are more than 20 dB below the impulse. Deviation from the triangular wave shape is due to bandwidth-limiting effects of input and output circuits. A spurious sidelobe -17 dB below the main impulse is caused by distortion present in the input waveforms. (b) Typical sidelobe structure of repeated code. Note change in time-scale.



E. CONCLUSIONS

All principal design goals for acoustoelectric convolvers have been met, and the preliminary results suggest that a convolver should function exceptionally well in systems with pseudorandom phase-shift-keyed modems. It is therefore recommended that acoustoelectric convolver subsystems be developed for packet radio repeaters.

F. COHERENT INTEGRATORS

The storage of a reference signal with an array of Schottky diodes, and the cross-correlation of subsequent signals with this reference was described in an article in Applied Physics Letters in June 1975.* We have modified this structure to provide for the coherent overlay of a succession of RF signals. In the aforementioned reference, Schottky diodes are contacted with metal overlays that act as capacitors, and an equivalent circuit of these diodes and capacitors is shown in Fig. II-11, in which R_s = Schottky diode resistance, C_s = Schottky diode capacity, and C_p = plate capacitor. When the Schottky diodes are forward biased $R_s^+ C_s^+ = 10^{-9}$ sec, and when reverse biased $R_s^- C_s^- = 10^{-1}$ sec. If the diodes are forward biased with a voltage impulse which is short in time relative to the RF period, then the piezoelectric field associated with an acoustic wave on the adjacent surface of LiNbO_3 causes a current pulse to flow from the diode contact into the silicon, which results in a charge distribution in the diode array which contains a replica of the acoustic-wave field. These charges, and the electrostatic fields associated with them, distort the underlying carrier distribution of the silicon. A subsequent surface-wave signal will cross-correlate with this carrier distribution, and a correlation signal appears at the output terminal.




Fig. II-11. Schematic showing a silicon strip with a Schottky diode matrix mounted $0.2 \mu\text{m}$ away from the LiNbO_3 surface-wave delay line. Also shown is a photograph of the Schottky diode matrix and an equivalent circuit of one diode.

Phase-coherent integration of a succession of reference signals requires the overlay of one signal on another. This is not possible with this Schottky diode circuit, because the previously stored signal is lost, within a nanosecond whenever the diodes are forward biased.

*K. A. Ingebrigtsen, R. A. Cohen, and R. W. Mountain, Appl. Phys. Lett. (June 1975).

We have modified the Schottky diode array by replacing the metallic overlay with a polycrystalline-silicon film overlay. This has the effect of putting in series with C_p a resistor R_p , with values consistent with $R_p C_p \sim 10^{-6}$ sec. In our device, C_p is approximately 16 times larger than C_s with zero or slightly reverse-biased diodes. Thus, most (16/17) of the charge on C_s is transferred to C_p in several microseconds. If the Schottky diodes are forward biased a second time, then an additional charge is transferred from C_s to C_p . If this is done with successive pulses which overlay exactly the previous pulses, then a large number of signals are coherently summed up to the point where the potential on C_p approaches the momentary potential on C_s .

The charge leaks to the bulk silicon through the back-biased diodes with a time constant $\tau = R_s(C_s + C_p)$, which is approximately 100 msec. A 2.75-cm \times 20-mm strip of n-type silicon with a bulk resistivity of 30 ohm-cm was used. The silicon interface is covered with a uniform array of platinum-silicide Schottky diodes each with a diameter of 5 μ m. The diodes are made by etching 5- μ m-diameter holes in a 0.1- μ m-thick oxide. These holes are overlaid with a thin layer of RF-sputtered platinum, and the wafer is heated to 600°C. Each diode is overlaid with a 10-mm \times 10- μ m square of polysilicon having a sheet resistance of 10^{10} ohms/square. The Schottky diode array is held 0.2 μ m away from the adjacent lithium niobate with etched spacer posts on the lithium niobate delay line.

Since the back-biased resistance of the Schottky diodes, and the resistance of the polysilicon is quite high, the piezoelectric field of the surface wave terminates on charges in the bulk silicon. A 500-V, 5-nsec voltage pulse is applied across the silicon-LiNbO₃ structure to forward bias the Schottky diodes a few tenths of a volt for a fraction of an RF period. This causes some of the piezoelectric field to terminate on the diode contacts, and a charge is transferred from the bulk silicon which is proportional to the instantaneous local electric displacement. Thus, a charge image of the acoustic wave field plus a uniform charge due to the DC pulse is obtained. This charge back-biases the diodes and the diode leakage current establishes a decay rate of about 0.1 sec.

Most of the 500-V forward-bias pulse appears across the 0.75-mm-thick lithium niobate crystal. A thinner crystal would require a proportionally lower voltage.

The plate output voltage in Fig. II-12 is caused by the correlation of a stored RF pulse at a frequency of 70 MHz and a time duration of 1.8 μ sec with a similar reading pulse. The reading pulse occurred 10 msec after the reference pulse was stored. The slight distortion of the triangle




Fig. II-12. Plate output signal resulting from the correlation of a stored RF pulse at 70 MHz and of time duration 1.8 μ sec with a similar reading pulse. Reading pulse occurred 10 msec after reference pulse was stored.

is caused by the attenuation of the acoustic signals. The efficiency of the device in terms of output power (P_{out}), stored reference signal input power (P_w), and reading pulse power (P_r) is

$$P_{out} - P_w - P_r = -83 \text{ dBm}$$

for a reference storage time of 1 msec. The correlation signal is linearly proportional to the reference signal up to a power level of 30 dBm and with reading signals up to 25 dBm.

The polycrystalline-silicon film of sheet resistivity 10^{10} ohms/square has associated with it a relaxation time of several microseconds. This leaves any charge accumulated in the polysilicon unaffected by subsequent 5-nsec-long high-voltage pulses. A charge brought to the platinum-silicide dot by a high-voltage pulse will require several microseconds to stream into the polysilicon island. Since the capacity C_p is substantially larger than the capacity C_s of the Schottky dot, most of the charge ends up in the polysilicon. If the high-voltage pulse is exactly timed with respect to the input signal, then the charge in the polysilicon should build up in phase with consecutive writings, and it should be possible to coherently integrate a signal.




Fig. II-13. Plate output signal during a coherent integration of the stored reference pattern. Horizontal scale is 20 msec/div. A burst of twelve pulses with an interpulse period of 8 msec occurs every 1.2 sec. Reading signals are free-running pulses with a repetition rate of 1.5×10^5 /sec. Output amplifier is blocked with a switch during the application of a 500-V, 5-nsec pulse which causes the dark vertical lines in the picture. The figure shows the buildup of the output signal during the initial integration period of 96 msec. The high-voltage pulses are then gated off to show the decay of the stored signal.

Figure II-13 shows the result of a coherent integration of an RF pulse. Twelve consecutive high-voltage pulses occur every 1.2 sec with a minimum interpulse period T_p of 8 msec. The output amplifier is blocked with a switch during application of the high-voltage pulse, which causes the dark vertical lines in the picture. A free-running reading pulse with a repetition rate of 1.5×10^5 /sec was used. The writing and reading pulses are identical to those used in connection with Fig. II-12. The figure shows a buildup of the output signal during the initial integration period of 96 msec. The high-voltage pulses are then gated off to show the decay of the stored signal.

The initial 20 msec of decay follows approximately an exponential law with a time constant of 30 msec, whereas the subsequent somewhat slower decay corresponds to a time constant of 60 msec. In this instance, the leveling of the output signal is related to the decay time.

In conclusion, we have demonstrated a novel analog signal processor with unique capabilities. It seems feasible to store with this device signals of 100-MHz bandwidth and 10- μ sec time duration, and to overlay coherently a succession of such signals over a time period of several tens of milliseconds, and then cross-correlate a reading signal with this accumulated sum.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER ESD-TR-76-101	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) Information Processing Techniques Program, Volume I: Packet Speech/Acoustic Convolvers		5. TYPE OF REPORT & PERIOD COVERED Semiannual Technical Summary 1 January - 30 June 1975	
6. AUTHOR(S) Bernard Gold Ernest Stern		7. PERFORMING ORG. REPORT NUMBER	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173		8. CONTRACT OR GRANT NUMBER(S) F19628-73-C-0002 ✓ ARPA Order-2006	
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order Nos. 2006 & 2929 Program Element Nos. 61101E & 62708E Project Nos. 6P10 & 6T10	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB Bedford, MA 01731		12. REPORT DATE 30 June 1975	
		13. NUMBER OF PAGES 32	
		15. SECURITY CLASS. (of this report) Unclassified	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		15a. DECLASSIFICATION DOWNGRADING SCHEDULE	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) [F]			
18. SUPPLEMENTARY NOTES None			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) information processing techniques packet speech network speech compression acoustic convolvers digital voice terminal ARPANET			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The Information Processing Techniques Program sponsored by DARPA at Lincoln Laboratory consists of three efforts: Packet Speech, Acoustic Convolvers, and Airborne Command and Control. In this Semiannual Technical Summary, the first two areas are reported in Vol. I and the third in Vol. II.			

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

207650 ✓

y/b